Spatially-constrained clustering of ecological networks

Vincent Miele*, Franck Picard and Stéphane Dray

Laboratoire Biométrie et Biologie Evolutive, Université de Lyon,
Université Lyon 1, CNRS, UMR5558; Villeurbanne, France.

arXiv:1402.5319v1 [stat.AP] 21 Feb 2014

Abstract

Spatial ecological networks are widely used to model interactions between georeferenced biological entities (e.g., populations or communities). The analysis of such data often leads to a two-step approach where groups containing similar biological entities are firstly identified and the spatial information is used afterwards to improve the ecological interpretation. We develop an integrative approach to retrieve groups of nodes that are geographically close and ecologically similar. Our model-based spatially-constrained method embeds the geographical information within a regularization framework by adding some constraints to the maximum likelihood estimation of parameters. A simulation study and the analysis of real data demonstrate that our approach is able to detect complex spatial patterns that are ecologically meaningful. The model-based framework allows us to consider external information (e.g., geographic proximities, covariates) in the analysis of ecological networks and appears to be an appealing alternative to consider such data.

Key-words: Graph Laplacian; Model-based clustering; Stochastic block model; Regularized EM-algorithm; Spatial partitioning; Spatial structure; Hydrothermal vents.

Introduction

In many ecological studies, researchers must analyze data describing the interactions between biological entities (e.g., individuals, populations, species or communities). These interactions can be directly observed (e.g., trophic relationships in a food-web, Krause et al., 2003) or they can be inferred from computed distance/similarity measures. Several genetic distances have been developed to summarize allele frequency differences between populations for instance (Kalinowski, 2002). In community ecology, species abundance (or presence/absence) data are routinely used to assess turnover (i.e., beta-diversity) between sites based on differences in species composition (Chapter 7 in Legendre and Legendre, 2012). Describing and summarizing these sets of pairwise interactions is an important step to better understand the functioning of ecological systems. In a theoretical viewpoint, networks (or graphs) offer a natural and efficient framework to store and analyze interaction data. Whereas food-web analysis was a traditional and historical field (Ings et al., 2009), network analysis techniques have recently gained popularity in genetics (Keller et al., 2013; Albert et al., 2013), movement ecology (Jacoby et al., 2012), landscape ecology (Bodin and Norberg, 2006; Pereira et al.,

*Corresponding author: vincent.miele@univ-lyon1.fr

2011), biogeography (Thébault, 2013; Moalic et al., 2012) or species distribution modeling (Foltête et al., 2012; Arajo et al., 2011).

From a statistical perspective, interaction data can be handled using *ecological networks* where biological entities correspond to the nodes while the intensities of interactions are represented by weighted edges (Proulx et al., 2005). Traditional approaches consist in summarizing the structure of a network using easy-to-compute statistics that relate to ecological properties (connectivity, degree distribution, average path length, Rayfield et al. (2011)). Unsupervised clustering is another common practice to detect modules *i.e.* groups of biological entities more densely connected to each other than to other external entities. In ecology, this strategy has been widely applied: several works have linked the modularity to the stability of the ecological network (Krause et al., 2003); finding modules has also helped to identify co-evolution patterns (Dupont and Olesen, 2009), to delineate conservation units (Fortuna et al., 2009), habitat patches (Pereira et al., 2011) or biogeographic entities (Moalic et al., 2012).

The identification of groups of nodes (such as modules for instance) has received considerable interest in physical sciences (see Newman (2006) or Fortunato (2010) for an extensive review) but also in statistics: in particular this work relies on well-established model-based clustering procedures (Daudin et al., 2008; Picard et al., 2009). These methods are fundamentally different from classic approaches by assuming a statistical distribution for interaction data. They are thus very flexible for handling various types of interaction data by specifying adequate statistical distribution to model presence/absence (Bernoulli), abundances (Poisson) or fluxes (Gaussian) or to integrate external covariates (Mariadassou et al., 2010). This general statistical framework is adequate for using tools from model selection theory (such as Bayesian Information Criterion (BIC)) in order to determine the number of groups that structure the data. Lastly, these methods can detect hidden connectivity patterns that are not limited to modularity, such as centrality (hubs) or hierarchy (see Picard et al. (2009) on a food web of wasps).

In many situations, the nodes correspond to entities that have explicit geographic locations transforming ecological networks into *spatial networks* (Dale and Fortin, 2010). When available, this spatial information is often used *posterior* to the identification of modules to improve their ecological interpretation (Moalic et al., 2012; Arajo et al., 2011; Dáttilo et al., 2013). However, if the aim of a study is to identify spatially-coherent modules (e.g., habitat patches), this indirect approach may not be optimal as it considers the spatial aspect only after summarizing the network structure. An appropriate technique would integrate spatial information explicitly in the detection of modules. This process is the core of spatially-constrained clustering techniques (Ambroise and Govaert, 1998; Gordon, 1996; Duque et al., 2007) when ecological observations are stored as raw data tables. When data consists in pairwise interactions stored as a network, no direct technique yet exists to delimit spatial modules except related strategies based on boundary detection (Monmonier, 1973).

In this work, we propose an original method to integrate the spatial information in the clustering of ecological interactions. Whereas spatial information is traditionally treated as geographic coordinates, trend surfaces or distance matrix, we used a more efficient strategy that models spatial proximity as a *structural network* where two nodes are connected if they are considered as neighbors (Dray et al., 2006, 2012). This approach offers a great flexibility allowing to include the effect of structural constraints (landscape fragmentation) or physical

barriers (rivers, mountains). In our framework, what is usually refered to a spatial network is thus decomposed into a pair of networks: nodes are identical but edges reflect either biological interactions (ecological network) or spatial connectivity (structural network). Our new model-based approach considers both networks in a single step allowing to identify spatially-coherent modules. The originality of our method is to embed the geographical information within a regularization framework (Ambroise and Govaert, 1998; He et al., 2011) that has been popularized by the analysis of high dimensional data sets (Hastie et al., 2001; Buhlmann and van de Geer, 2011). We evaluate our new approach by a simulation study and applied it on a real ecological data set. An R package including methods and data to perform the analysis is also provided at `http://lbbe.univ-lyon1.fr/geoclust`.

Materials and Methods

Model-based Clustering of Interaction Data

Our method belongs to the general framework of model-based clustering of network data. This family of models includes the stochastic block model (Airoldi et al., 2008) and the MixNet approach (Daudin et al., 2008; Mariadassou et al., 2010) among others. We consider data for n interacting entities with Y_{ij} standing for the observed measure of these interactions and $\mathbf{Y} = (Y_{ij})$. Our model is based on group-membership of entities and $\mathbf{Z} = (\mathbf{Z}_1, ..., \mathbf{Z}_n)$ denotes the matrix of labels of entities $1, .., n$, *i.e.* $Z_{iq} = 1$ if i belongs to group q and 0 otherwise. In the context of unsupervised clustering, this matrix is unknown and our method aims to recover these labels using the observed information contained in $\mathbf{Y}$. Model-based clustering hypotheses that if labels were known, the distribution of the interaction data would be completely determined. Hence, we start by assuming that there are Q groups with proportions $\boldsymbol{\alpha} = (\alpha_1, ..., \alpha_Q)$ such that the distribution of labels $\mathbf{Z}_i = (Z_{i1}, ...Z_{iQ})$ is Multinomial with parameter $\boldsymbol{\alpha}$. The number of groups Q is unknown and will be estimated afterwards. The distribution of the interaction data is specified conditionally to the labels:

$$\mathbf{Z}_i \sim \mathcal{M}(1, \boldsymbol{\alpha}), \;\; \mathbf{Z}_j \sim \mathcal{M}(1, \boldsymbol{\alpha}), \;\; Y_{ij} | \{Z_{iq} Z_{i\ell} = 1\} \sim f(\cdot, \theta_{ql}), \tag{1}$$

where distribution $f(\cdot, \theta_{ql})$ can be Bernoulli to model presence-absence data (Airoldi et al., 2008; Daudin et al., 2008), Gaussian or Poisson to model fluxes or abundance data (Mariadassou et al., 2010). The parameters of this model are the proportions for the groups ($\boldsymbol{\alpha}$) and the parameters governing the conditional distribution of the observations ($\boldsymbol{\theta} = (\theta_{q\ell})$). In the following, we note $\boldsymbol{\gamma} = (\boldsymbol{\alpha}, \boldsymbol{\theta})$.

The objective is to estimate $\boldsymbol{\gamma}$ and to recover the unobserved labels of the data using the posterior expectation of membership $\mathbb{E}(\mathbf{Z}_i | \mathbf{Y})$. This is achieved using the EM-algorithm to maximize the observed-data likelihood denoted by $\log \mathcal{L}(\mathbf{Y}; \boldsymbol{\gamma})$. Unfortunately, the direct maximization of this likelihood is untractable due to the total number of possible partitions ($\mathcal{L}(\mathbf{Y}; \boldsymbol{\gamma}) = \prod_{\mathbf{Z}} \mathcal{L}(\mathbf{Y}, \mathbf{Z}; \boldsymbol{\gamma})$). Hence, we use an iterative algorithm that maximizes the

complete-data likelihood:

$$\begin{aligned}\log \mathcal{L}(\mathbf{Y}, \mathbf{Z}; \boldsymbol{\gamma}) &= \log \mathcal{L}(\mathbf{Z}; \boldsymbol{\alpha}) + \log \mathcal{L}(\mathbf{Y}|\mathbf{Z}; \boldsymbol{\theta}) \\ &= \sum_{iq} Z_{iq} \log(\alpha_q) + \sum_{ij,q\ell} Z_{iq} Z_{j\ell} \log f(Y_{ij}; \theta_{q\ell}) \end{aligned} \quad (2)$$

The labels being unknown, the algorithm proceeds as follows: the E-Step computes the conditional expectation of the complete-data log-likelihood defined as:

$$\begin{aligned}\mathcal{Q}(\boldsymbol{\gamma}, \boldsymbol{\gamma}^{[h]}) &= \mathbb{E}_{\boldsymbol{\gamma}^{[h]}} \{\log \mathcal{L}(\mathbf{Y}, \mathbf{Z}; \boldsymbol{\gamma}) | \mathbf{Y}\} \\ &= \sum_{iq} \mathbb{E}_{\boldsymbol{\gamma}^{[h]}}(Z_{iq}|\mathbf{Y}) \log(\alpha_q) + \sum_{ij,q\ell} \mathbb{E}_{\boldsymbol{\gamma}^{[h]}}(Z_{iq} Z_{j\ell}|\mathbf{Y}) \log f(Y_{ij}; \theta_{q\ell}), \end{aligned} \quad (3)$$

for a current value of the parameters ($\boldsymbol{\gamma}^{[h]}$). Then the M-step maximizes $\mathcal{Q}$ with respect to $\boldsymbol{\alpha}$ and $\boldsymbol{\theta}$. Computational difficulties often arise at the E-step, mainly due to complex dependency structures that can govern the *posterior* distribution of labels given the data. This issue has motivated many methodological developments, in particular in the context of network data, with the use of variational methods (Jordan et al., 1999) to compute this *posterior* distribution (Daudin et al., 2008; Mariadassou et al., 2010).

Accounting for Spatial Constraints in the Clustering Model

Labels regularization using a spatial network

We choose to use a structural network that records the spatial proximity between ecological entities of the ecological network $\mathbf{Y}$ (see Figure 1). Structural networks are sometimes directly available such as road networks (Smaltschinski et al., 2012), but they are usually constructed using geographical data with *ad-hoc* techniques such as maximum spanning trees (Assunção et al., 2006), k-nearest neighbors (Guo, 2008), distance threshodling (Pereira et al., 2011) or edge-thinning (Urban and Keitt, 2001; Keller et al., 2013). In the following we suppose that the structural network is given and fixed. It is denoted by $\mathbf{X} = (X_{ij})$ such that $(X_{ij} > 0)$ is the geographical proximity between entities i and j and $X_{ij} = 0$ if they are not connected. The entities are the same as in the ecological network $\mathbf{Y}$.

We propose to embed the geographical information within a regularization framework by adding some constraints in the maximum likelihood estimation of parameters. In regularization techniques, a constraint defined by a network can be introduced using the graph Laplacian (Jacob et al., 2012). For a network with connection matrix $\mathbf{X} = (X_{ij})$, the Laplacian is defined by $\mathbf{L_X} = \mathbf{D} - \mathbf{X}$ where $\mathbf{D}$ is the diagonal matrix of degrees with diagonal terms $d_i = \sum_j X_{ij}$. The Laplacian $\mathbf{L_X}$ can then be used as a metric to measure the spatial variability. Indeed, for a given vector $\mathbf{u} = (u_1, ...u_n)$, we have:

$$\|\mathbf{u}\|^2_{\mathbf{L_X}} = \mathbf{u}^T \mathbf{L_X} \mathbf{u} = \sum_{ij} X_{ij} (u_i - u_j)^2,$$

which is the squared distance between values of $\mathbf{u}$ weighted by their spatial proximities contained in $\mathbf{X}$. This quantity has been defined as the local variance by Lebart (1969) and is equal to the numerator of the spatial autocorrelation index proposed by Geary (1954).

We develop an original regularization procedure aiming to reduce the variation of labels along the spatial network. Whereas the vector of parameters is traditionally regularized, our approach considers that the vector of labels can be regularized using the spatial network $\mathbf{X}$. Denoting by $\mathbf{Z}^q = (Z_{1q}, \ldots, Z_{nq})$ the vector of individuals for label q, we propose the following penalty:

$$\text{pen}(\mathbf{Z}; \mathbf{L}_\mathbf{X}) = \sum_{q=1}^{Q} \|\mathbf{Z}^q\|^2_{\mathbf{L}_\mathbf{X}} = \sum_{q=1}^{Q} \sum_{i,j} X_{ij}(Z_{iq} - Z_{jq})^2,$$

and the likelihood to maximize by the EM-algorithm becomes:

$$\log \mathcal{L}(\mathbf{Y}, \mathbf{Z}; \boldsymbol{\gamma}) - \lambda \times \text{pen}(\mathbf{Z}; \mathbf{L}_\mathbf{X}),$$

with λ a constant controlling the amount of penalization that can be estimated adaptively to the data. Let us consider the case where $X_{ij} \in \{0, 1\}$ to interpret the penalty. In this case,

$$\text{pen}(\mathbf{Z}; \mathbf{L}_\mathbf{X}) = \sum_{q=1}^{Q} \sum_{i \sim j} (Z_{iq} - Z_{jq})^2 = \sum_{q=1}^{Q} \sum_{i \sim j} 1_{\{Z_{iq} \neq Z_{jq}\}}$$

with $i \sim j$ standing for entities i and j connected in the spatial network, so that the penalty accounts for the number of edges in the spatial network that have discordant labels.

Regularized EM-algorithm based on spatial network

Considering the new penalized likelihood, the regularized EM algorithm is based on the conditional expectation of the penalized complete-data likelihood

$$\mathcal{Q}(\boldsymbol{\gamma}, \boldsymbol{\gamma}^{[h]}) - \lambda \times \mathbb{E}_{\boldsymbol{\gamma}^{[h]}} \left\{ \text{pen}(\mathbf{Z}; \mathbf{L}_\mathbf{X}) | \mathbf{Y} \right\}.$$

As the penalty term does not involve the parameters but only the labels, the maximization step is unchanged (Mariadassou et al., 2010). For instance, if the interaction data are supposed to be Gaussian such that

$$Y_{ij} | \{Z_{iq} Z_{i\ell} = 1\} \sim \mathcal{N}(\mu_{q\ell}, \sigma^2), \tag{4}$$

we get the following updated estimates (at iteration $[h+1]$):

$$\widehat{\mu}_{ql}^{[h+1]} = \frac{\sum_{ij,q\ell} \widehat{Z}_{iq}^{[h]} \widehat{Z}_{j\ell}^{[h]} Y_{ij}}{\sum_{ij,q\ell} \widehat{Z}_{iq}^{[h]} \widehat{Z}_{j\ell}^{[h]}}, \quad \widehat{\sigma}^{2[h+1]} = \frac{\sum_{ij,q\ell} \widehat{Z}_{iq}^{[h]} \widehat{Z}_{jl}^{[h]} (Y_{ij} - \widehat{\mu}_{ql}^{[h+1]})^2}{\sum_{ij,q\ell} \widehat{Z}_{iq}^{[h]} \widehat{Z}_{j\ell}^{[h]}},$$

with $(\widehat{Z}_{iq}^{[h]})$ the predicted labels provided by the E-step (below).

Network-Based Regularized E-Step. The traditional EM algorithm is not tractable in the case of mixtures for interaction data (Daudin et al., 2008). Thus, we use a variational approach which consists in choosing a surrogate *posterior* distribution for the labels so that we can suppose that they are conditionally independent (Jordan et al., 1999). Denoting by

$\tau_{iq} \simeq \mathbb{E}_{\boldsymbol{\gamma}^{[h]}}\{Z_{iq}|\mathbf{Y}\}$ the approximate posterior expectation of labels and by $\boldsymbol{\tau} = (\boldsymbol{\tau}_i) = (\tau_{iq})$, the penalized complete-data likelihood becomes:

$$\begin{aligned}\mathcal{Q}(\boldsymbol{\gamma}, \boldsymbol{\gamma}^{[h]}) - \lambda \times \mathbb{E}_{\boldsymbol{\gamma}^{[h]}}\{\text{pen}(\mathbf{Z}; \mathbf{L}_{\mathbf{X}})|\mathbf{Y}\} &\simeq \sum_{iq} \tau_{iq} \log(\alpha_q) + \sum_{ij,q\ell} \tau_{iq}\tau_{j\ell} \log f(Y_{ij}; \theta_{q\ell}) \\ &- \lambda \times \sum_{i,j} X_{ij}(\boldsymbol{\tau}_i - \boldsymbol{\tau}_j)^2.\end{aligned}$$

The computation of these approximate posterior probabilities is achieved by solving a fixed-point algorithm (Ambroise and Govaert, 1998; Mariadassou et al., 2010) and we account for spatial constraints through the penalty term. The derivation of this algorithm is provided in the Appendix. Lastly, to retrieve clear and separable groups, we add a classification step corresponding to the Classification-EM algorithm (Celeux and Govaert, 1992). We use the Maximum *a posteriori* rule such that $\widehat{Z}_{iq} = 1$ if $q = \arg\max_\ell \widehat{\tau}_{i\ell}$ and 0 otherwise.

Initialization, λ values and number of groups It is well know that EM-algorithms are very sensitive to the quality of the initialization point. We propose to set-up $\boldsymbol{\tau}^{(0)}$ using the clustering partition $\widehat{\mathbf{Z}}^0$ obtained with a traditional k-means algorithm as suggested by Daudin et al. (2008). Parameter λ increased until $\widehat{\lambda}_{\max}$ that corresponds to the maximal spatial homogeneity (i.e. until no further change of labels). Lastly, the number of groups $\widehat{Q}(\widehat{\lambda}_{\max})$ is chosen using model selection strategy based on the Integrated Classification Likelihood (ICL, Daudin et al. (2008)).

Results

Simulation study

In this section we evaluate the relative merits of our method on simulated data that explore different degrees of coherence between the interaction data and the spatial information. For standard clustering methods, the performance only depends on the *group separability*, *i.e.* the ability to delineate groups within the ecological network. However, in the case of spatially constrained clustering, performance is more difficult to explore since it depends on a trade-off between group separability and spatial homogeneity/coherence of the groups. The easiest configuration is when the groups and the spatial information overlap (low *spatial discordance*), the most difficult situation is when groups and spatial information are contradictory (high *spatial discordance*).

We start by simulating a structural network $\mathbf{X}$ that consists of two spatial components of 50 nodes each, related by one single edge only (see Supplementary Figure 1). The proximity values stored in the edges in $\mathbf{X}$ are determined using a Gabriel graph (Dale and Fortin, 2010). Then we consider $Q = 2$ groups with equal proportions $\boldsymbol{\alpha}^* = (0.5, 0.5)$, and we sample the true labels $\mathbf{Z}^*$ such that $\mathbf{Z}_i \sim \mathcal{M}(1, \boldsymbol{\alpha}^*)$. In order to explore different degrees of spatial discordance for the groups, we start by matching labels with the connected components of $\mathbf{X}$ (no spatial discordance), and we randomly sample an increasing number of pairs of nodes from the two spatial components and swap their labels (without modifying $\mathbf{X}$). Then we use the proportion of edges in $\mathbf{X}$ with discordant labels in the spatial network as a measure of

spatial discordance. In the most favorable situation, there is a perfect overlap between labels and the spatial components in $\mathbf{X}$. In the worst situation, labels are spread independently of $\mathbf{X}$. Once the spatial discordance of labels is determined, we sample a Gaussian interaction network $\mathbf{Y}$ (Eq. 4) with parameters $\sigma = 0.2$ and mean parameters:

$$\boldsymbol{\mu} = \begin{bmatrix} \mu & \nu \\ \nu & \mu \end{bmatrix}.$$

The group separability is governed by $\delta = (\mu - \nu)/\sigma$ which varied from no signal ($\delta = 0$) to high signal levels ($\delta = 1$). Each configuration was repeated 50 times and we applied our method with spatial penalty or without (*i.e.* $\lambda = 0$). The performance of the method is then assessed by the averaged adjusted Rand Index (aRI, Hubert and Arabie (1985)) that measures the agreement between two partitions. It lies between 0 and 1, 1 being for a perfect match between the simulated and the recovered partition. This index is first used to assess the classification of interactions (between estimated and simulated labels $\mathbf{Z}^*$, referred to as *aRI on* $\mathbf{Y}$, with the true classification given by the two spatial components), and secondly to assess the spatial homogeneity of groups (between estimated labels and the two spatial components of $\mathbf{X}$, referred to as *aRI on* $\mathbf{X}$).

When groups are highly separable ($\delta \geq 0.5$) and without any spatial discordance, clustering performs perfectly with or without the spatial penalty (aRI $\simeq 1$ on both $\mathbf{X}$ and $\mathbf{Y}$, Figure 2A-D). The interest in using the spatial penalty is illustrated when the spatial discordance increases until the proportion of spatially discordant labels reaches 50%. In this case, the spatial penalty induces a lower classification performance on the interaction data (decreasing aRI on $\mathbf{Y}$, Figure 2B), which is compensated by an excellent clustering performance on spatial information (aRI>0.9 on $\mathbf{X}$, Figure 2A). Consequently, the effect of the spatial penalty is to maintain the spatial homogeneity of the groups at the price of classification errors on $\mathbf{Y}$. When the separability of the group decreased ($\delta = 0.25$), the performance of clustering also decreased, with a aRI of 0.3 on $\mathbf{Y}$. In this case, the spatial penalty clearly helps to recover the groups when they match the spatial structure (aRI=0.8 on $\mathbf{X}$ without spatial discordance, $\delta = 0.25$, Figure 2A,C). In conclusion, when the spatial discordance is low to moderate, the spatial penalty helps in finding the groups, even if the signal in the interaction network is moderate.

Lastly, we evaluate the ability of our method to identify groups when the spatial discordance is high. When looking at the number of groups chosen by the model selection procedure (Figure 3), it is clear that all nodes are gathered in a single group when the spatial discordance increases. This proves that our method is adaptive to the spatial discordance, and searches only for *spatially-coherent modules*: when there is no spatial signal, our method returns no clustering instead to identify groups that are not spatially organized.

Hydrothermal vents

We re-examine data concerning the presence/absence of 332 genera in 63 hydrothermal fields. This data set was used by Bachraty et al. (2009) and Moalic et al. (2012) to define biogeographic provinces based on the faunal distribution in oceanic hydrothermal vents. We built an ecological network by computing distances in faunal composition using Jaccard coefficient

(see Chapter 7 in Legendre and Legendre, 2012). This fully connected network contains 63 nodes (the fields) and edges are weighted by the ecological distances.

The distribution of Jaccard coefficients was bimodal due to an excess of distance values equal to 1 corresponding to pairs of fields sharing no species; the remaining values, after a logistic transform, follow a Gaussian distribution. Consequently, we modelled data using an inflated Gaussian distribution. We computed spatial proximities among the fields by computing great circle distances. The structural network was then built by removing all the edges corresponding to distances higher than 3600 km (see Figure 5B). This choice ensures that each node has at least one edge and that only local spatial proximities are considered. Edges were then weighted by spatial distances ($max(d_{ij}) - d_{ij}$).

We first applied our algorithm without spatial penalty (i.e., $\lambda = 0$). Fields were partitioned in 6 groups based on their species composition (Figure 4). It is clear that fields share more ecological similarities within groups that between groups (comparison of diagonal and off-diagonal terms respectively in Figure 4A). This partitioning highlights some geographical patterns: some groups are roughly spatially coherent (groups A, C, D, E) whereas other groups did not show any spatial coherence (Figure 4B).

If the spatial penalty is applied, the model selection procedure selected also a clustering in 6 groups (Figure 5B). In this case, within-group similarities are also higher than the between-group (Figure 5A). Average within-group Jaccard index (0.944) is slightly higher than for the unconstrained approach (0.928). This loss of ecological homogeneity is counterbalanced by a gain in the spatial aspect (Figure 5B) as all ecological groups are now spatially coherent. Our results are very similar to those obtained by Bachraty et al. (2009) and Moalic et al. (2012) who used different methods and identified 6 and 5 provinces, respectively. Their partitioning disagreed for the provinces in West Pacific (WP), Indian Ocean (IO) and East Pacific Rise (EPR). We identified a group with only one field (Loihi Seamount, LOS) that share no species with all others and is poorly spatially connected (only one edge) in the structural network. As in the mentioned studies, we found the Mid-Atlantic Ridge (MAR) and the NorthEast Pacific (NE) provinces, the latter including the Guaymas Basin in our results. We also identified the partitioning between the Northern and Southern-Pacific Rise provinces (NEPR and SEPR) that was proposed in Bachraty et al. (2009) as centers of dispersal for hydrothermal fauna. This shows that our algorithm can separate groups of fields that are spatially connected but very dissimilar according to the ecological network. Our algorithm was not able to partition the global province (NW+SW+IO on Figure 5B) composed by the Indian Ocean (IO) and Western Pacific zones. Interestingly, both previous studies were discordant for this region: Bachraty et al. (2009) merged Indian Ocean and Southwest Pacific (SW) and introduced an additional Northwest Pacific province (NW) whereas Moalic et al. (2012) proposed a global WP province but kept aside the IO province. Moalic et al. (2012) focused on the central role of the WP province in the biogeography of hydrothermal vents and also mentioned that the existence of the IO province remains unclear due its poor representation in the data. Consequently, the group NW+SW+IO is the most ecologically heterogeneous in our results (Figure 5A) because our method is more conservative and suggests that any partitioning would be artificial in this region. This merging of IO and Western Pacific zones (SW, NW) demonstrated that our algorithm is able to merge two groups that are ecologically similar but not spatially connected.

Conclusions

The analysis of ecological networks has gained a lot of attention in the past years, and the development of dedicated statistical methods has been challenging. Here we propose a statistical framework to identify spatial groups of ecological entities based on interaction data. We developed an efficient way to retrieve spatially-coherent groups of ecological entities using a regularization framework. The analysis of simulated data and of a real dataset demonstrated the ability of our method to identify groups that match the spatial structure. If no spatial patterns is present in the data, our method did not identify any partitioning and standard clustering algorithms, which do not consider the spatial aspect, should be preferred in these circumstances.

Whereas classical analysis consists in describing the main structures of ecological networks by summary statistics, our method is based on a model-based framework that assumes and estimates a distribution for interaction data. This approach is already used to analyze genomics data but it has been rarely applied to ecological data (but see Picard et al., 2009). We demonstrated that this framework is very promising as it allows to include easily spatial constraints using a global model for the full network. It can be extended to build more complex models by considering covariates measured either on the nodes or the edges of the network (e.g., environmental variables). Lastly, different distributions can be modeled so that the method is applicable to a wide-variety of interaction data. For instance, we could assume a multivariate distribution to handle multi-layer ecological networks that record different types of interaction between entities (Kéfi et al., 2012).

Acknowledgements

The authors would like to thank Sophie Arnaud-Haond and Pierre Legendre for providing access to the hydrothermal fields dataset.

Appendix

The fixed point algorithm arises when considering:

$$\begin{aligned}\frac{\partial}{\partial \tau_{iq}}\left(\mathcal{Q}(\boldsymbol{\gamma},\boldsymbol{\gamma}^{[h]}) - \lambda \times \mathbb{E}_{\boldsymbol{\gamma}^{[h]}}\left\{\text{pen}(\mathbf{Z};\mathbf{L}_{\mathbf{X}})|\mathbf{Y}\right\}\right) &= -\log\ \tau_{iq} - 1 + \log \alpha_q + \sum_{ji,\ell} \tau_{j\ell}\ \log f(Y_{ij};\theta_{q\ell}) \\ &- 2\lambda \sum_{ij} X_{ij}(\tau_{iq} - \tau_{jq}) + L_i,\end{aligned}$$

where L_i is the Lagrange multiplier that ensures the constraint $\sum_q \tau_{iq} = 1$. Then the optimal parameter $\widehat{\tau}_{iq}$ satisfies:

$$\widehat{\tau}_{iq} \propto \alpha_q \left(\prod_{j\neq i}\prod_{\ell} f(Y_{ij};\theta_{q\ell})^{\widehat{\tau}_{jl}}\right) \exp\left(-2\lambda \sum_{i,j} X_{ij}(\widehat{\tau}_{iq} - \widehat{\tau}_{jq})\right)$$

References

Airoldi, E., Blei, D., Fienberg, S., and Xing, E. (2008). Mixed-membership stochastic blockmodels. *Journal of Machine Learning Research*, 9:1981–2014.

Albert, E. M., Fortuna, M. A., Godoy, J. A., and Bascompte, J. (2013). Assessing the robustness of networks of spatial genetic variation. *Ecology Letters*, 16 Suppl 1:86–93.

Ambroise, C. and Govaert, G. (1998). Convergence of an EM-type algorithm for spatial clustering. *Pattern Recognition Letters*, 19(10):919–927.

Arajo, M. B., Rozenfeld, A., Rahbek, C., and Marquet, P. A. (2011). Using species co-occurrence networks to assess the impacts of climate change. *Ecography*, 34(6):897–908.

Assunção, M., Corrêa Neves, M., Câmara, G., and da Costa Freitas, C. (2006). Efficient regionalization techniques for socio-economic geographical units using minimum spanning trees. *International Journal of Geographical Information Science*, 20(7):797–811.

Bachraty, C., Legendre, P., and Desbruyères, D. (2009). Biogeographic relationships among deep-sea hydrothermal vent faunas at global scale. *Deep Sea Research Part I: Oceanographic Research Papers*, 56(8):1371–1378.

Bodin, O. and Norberg, J. (2006). A network approach for analyzing spatially structured populations in fragmented landscape. *Landscape Ecology*, 22(1):31–44.

Buhlmann, P. and van de Geer, S. (2011). *Statistics for High-Dimensional Data: Methods, Theory and Applications*. New York: Springer-Verlag.

Celeux, G. and Govaert, G. (1992). A classification EM algorithm for clustering and two stochastic versions. *Computational Statistics and Data Analysis*, 14:315–332.

Dale, M. and Fortin, M.-J. (2010). From graphs to spatial graphs. *Annual Review of Ecology, Evolution, and Systematics*, 41(1):21–38.

Dáttilo, W., Guimarães, P. R., and Izzo, T. J. (2013). Spatial structure of ant-plant mutualistic networks. *Oikos*, 122(11):1643–1648.

Daudin, J. J., Picard, F., and Robin, S. (2008). A mixture model for random graphs. *Statistics and Computing*, 18(2):173–183.

Dray, S., Legendre, P., and Peres-Neto, P. R. (2006). Spatial modeling: a comprehensive framework for principal coordinate analysis of neighbor matrices (PCNM). *Ecological Modelling*, 196:483–493.

Dray, S., Pélissier, R., Couteron, P., Fortin, M. J., Legendre, P., Peres-Neto, P. R., Bellier, E., Bivand, R., Blanchet, F. G., De Caceres, M., Dufour, A. B., Heegaard, E., Jombart, T., Munoz, F., Oksanen, J., Thioulouse, J., and Wagner, H. H. (2012). Community ecology in the age of multivariate multiscale spatial analysis. *Ecological Monographs*, 82(3):257–275.

Dupont, Y. L. and Olesen, J. M. (2009). Ecological modules and roles of species in heathland plant-insect flower visitor networks. *Journal of Animal Ecology*, 78(2):346–353.

Duque, J. C., Ramos, R., and Surinach, J. (2007). Supervised regionalization methods: a survey. *International Regional Science Review*, 30:195–220.

Foltête, J.-C., Clauzel, C., Vuidel, G., and Tournant, P. (2012). Integrating graph-based connectivity metrics into species distribution models. *Landscape Ecology*, 27(4):557–569.

Fortuna, M. A., Albaladejo, R. G., Fernández, L., Aparicio, A., and Bascompte, J. (2009). Networks of spatial genetic variation across species. *Proceedings of the National Academy of Sciences of the United States of America*, 106(45):19044–19049.

Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, 486:75–174.

Geary, R. C. (1954). The contiguity ratio and statistical mapping. *The Incorporated Statistician*, 5(3):115–145.

Gordon, A. (1996). A survey of constrained classification. *Computational Statistics & Data Analysis*, 21(1):17–29.

Guo, D. (2008). Regionalization with dynamically constrained agglomerative clustering and partitioning (redcap). *International Journal of Geographical Information Science*, 22(7):801–823.

Hastie, T., Tibshirani, R., and Friedman, J. H. (2001). *The elements of statistical learning: data mining, inference, and prediction: with 200 full-color illustrations*. New York: Springer-Verlag.

He, X., Cai, D., Shao, Y., Bao, H., and Han, J. (2011). Laplacian regularized gaussian mixture model for data clustering. *IEEE Knowledge and Data Engineering*, pages 1406–1418.

Hubert, L. and Arabie, P. (1985). Comparing partitions. *Journal of Classification*, 2(1):193–218.

Ings, T. C., Montoya, J. M., Bascompte, J., Blüthgen, N., Brown, L., Dormann, C. F., Edwards, F., Figueroa, D., Jacob, U., Jones, J. I., Lauridsen, R. B., Ledger, M. E., Lewis, H. M., Olesen, J. M., van Veen, F. J. F., Warren, P. H., and Woodward, G. (2009). Ecological networks - beyond food webs. *Journal of Animal Ecology*, 78:253–269.

Jacob, L., Neuvial, P., and Dudoit, S. (2012). More power via graph-structured tests for differential expression of gene networks. *The Annals of Applied Statistics*, 6(2):561–600.

Jacoby, D. M. P., Brooks, E. J., Croft, D. P., and Sims, D. W. (2012). Developing a deeper understanding of animal movements and spatial dynamics through novel application of network analyses. *Methods in Ecology and Evolution*, 3:574583.

Jordan, M., Ghahramani, Z., Jaakkola, T., and Saul, L. (1999). An introduction to variational methods for graphical models. *Mach. Learn.*, 37(2):183–233.

Kalinowski, S. T. (2002). Evolutionary and statistical properties of three genetic distances. *Molecular Ecology*, 11:1263–1273.

Kéfi, S., Berlow, E. L., Wieters, E. A., Navarrete, S. A., Petchey, O. L., Wood, S. A., Boit, A., Joppa, L. N., Lafferty, K. D., Williams, R. J., Martinez, N. D., Menge, B. A., Blanchette, C. A., Iles, A. C., and Brose, U. (2012). More than a meal integrating non-feeding interactions into food webs. *Ecology Letters*, 15(4):291–300.

Keller, D., Holderegger, R., and van Strien, M. J. (2013). Spatial scale affects landscape genetic analysis of a wetland grasshopper. *Molecular Ecology*, 22(9):2467–2482.

Krause, A. E., Frank, K. A., Mason, D. M., Ulanowicz, R. E., and Taylor, W. W. (2003). Compartments revealed in food-web structure. *Nature*, 426(6964):282–285.

Lebart, L. (1969). Analyse statistique de la contiguïté. *Publication de l'Institut de Statistiques de l'Université de Paris*, 28:81–112.

Legendre, P. and Legendre, L. (2012). *Numerical Ecology*. Elsevier, Amsterdam, third edition.

Mariadassou, M., Robin, S., and Vacher, C. (2010). Uncovering latent structure in valued graphs: A variational approach. *Annals of Applied Statistics*, 4:715–742.

Moalic, Y., Desbruyères, D., Duarte, C. M., Rozenfeld, A. F., Bachraty, C., and Arnaud-Haond, S. (2012). Biogeography revisited with network theory: Retracing the history of hydrothermal vent communities. *Systematic Biology*, 61(1):127–137.

Monmonier, M. (1973). Maximum-difference barriers: an alternative numerical regionalization method. *Geographical Analysis*, 5(3):245–261.

Newman, M. E. J. (2006). Modularity and community structure in networks. *Proceedings of the National Academy of Sciences of the United States of America*, 103(23):8577–8582.

Pereira, M., Segurado, P., and Neves, N. (2011). Using spatial network structure in landscape management and planning: A case study with pond turtles. *Landscape and Urban Planning*, 100(1-2):67–76.

Picard, F., Miele, V., Daudin, J. J., Cottret, L., and Robin, S. (2009). Deciphering the connectivity structure of biological networks using MixNet. *BMC Bioinformatics*, 10 Suppl 6:S17.

Proulx, S. R., Promislow, D. E. L., and Phillips, P. C. (2005). Network thinking in ecology and evolution. *Trends in Ecology and Evolution*, 20(6):345–353.

Rayfield, B., Fortin, M.-J., and Fall, A. (2011). Connectivity for conservation: a framework to classify network measures. *Ecology*, 92(4):847–858.

Smaltschinski, T., Seeling, U., and Becker, G. (2012). Clustering forest harvest stands on spatial networks for optimised harvest scheduling. *Annals of Forest Science*, 69(5):651–657.

Thébault, E. (2013). Identifying compartments in presence-absence matrices and bipartite networks: insights into modularity measures. *Journal of Biogeography*, 40(4):759–768.

Urban, D. and Keitt, T. (2001). Landscape connectivity: a graph-theoretic perspective. *Ecology*, 82(5):1205–1218.

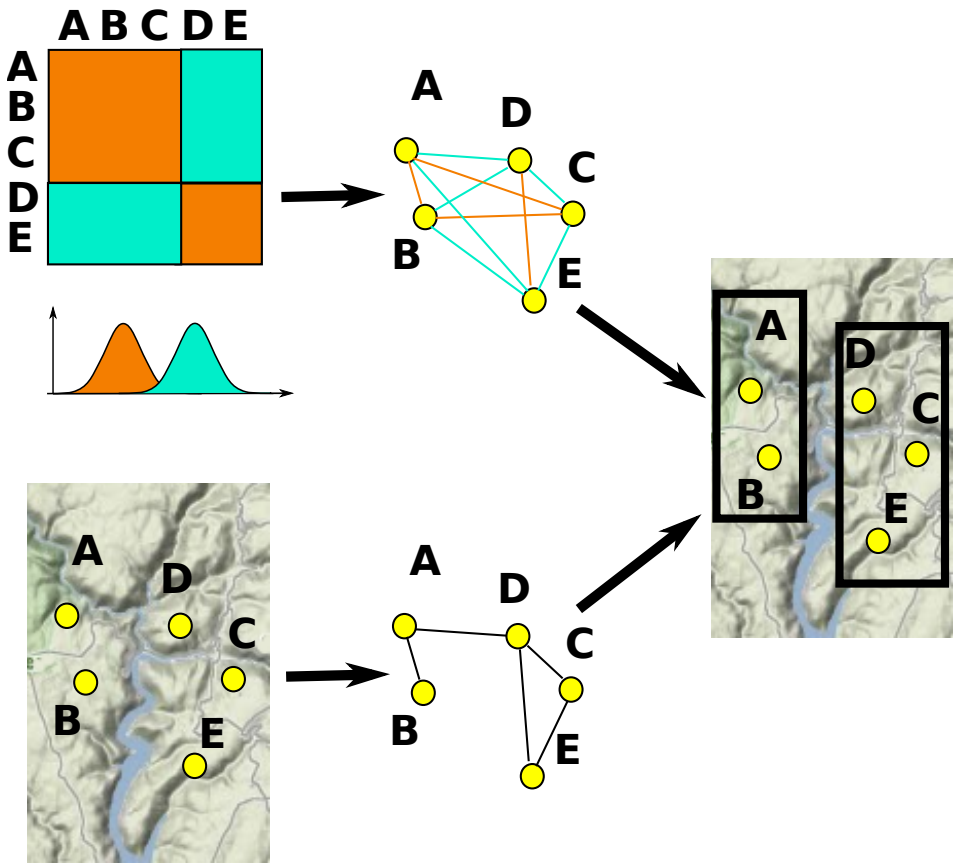

Figure 1: Framework of our algorithm. The ecological network (top) records the ecological distance between entities. The structural network (bottom) summarizes the proximity between geographical locations. Our method deciphers groups of entities (black squares on the right) using both networks into a model-based strategy associated to a regularization framework. While entities A,B,C and D,E forms two groups in the ecological data, the geographical constraints leads to lastly grouping C with D and E.

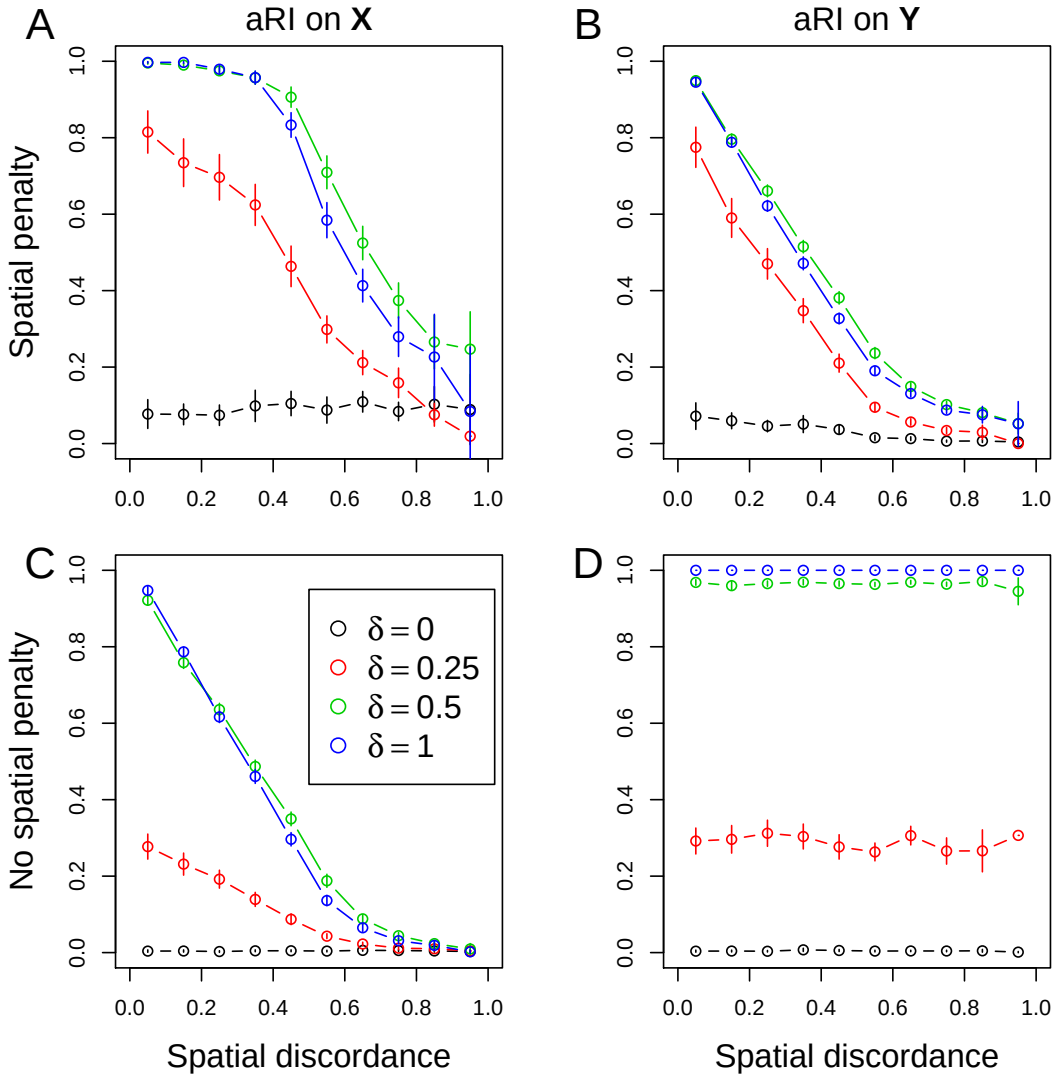

Figure 2: Performance of our method on simulated data, assessed by the averaged Adjusted Rand Index (aRI), for varying spatial discordance and group separability (δ) values. (A) aRI on $\mathbf{X}$ with our regularized EM algorithm with spatial penalty. (B) same for aRI on $\mathbf{Y}$. (C) aRI on $\mathbf{X}$ without spatial penalty, *i.e* $\lambda = 0$. (D) same for aRI on $\mathbf{Y}$.

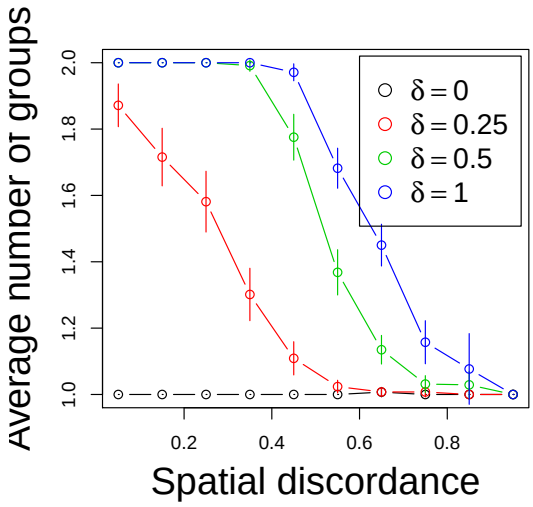

Figure 3: Average number of groups selected by our model selection approach, on simulated data, for varying spatial discordance and group separability (δ) values.

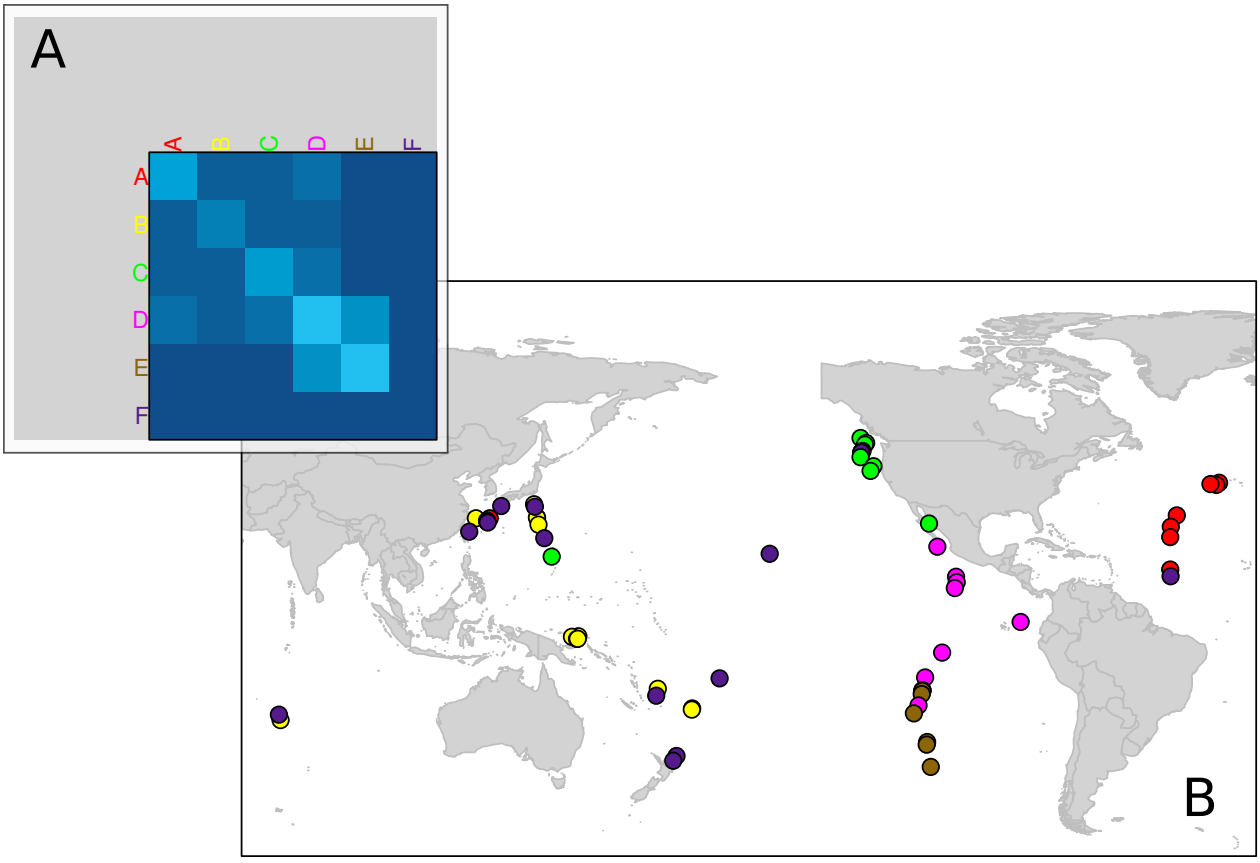

Figure 4: Clustering of 63 hydrothermal fields in 6 groups based on similarities in species composition (no spatial penalty). (A) Heatmap showing average ecological distances within and between groups. (B) Map showing the 6 groups

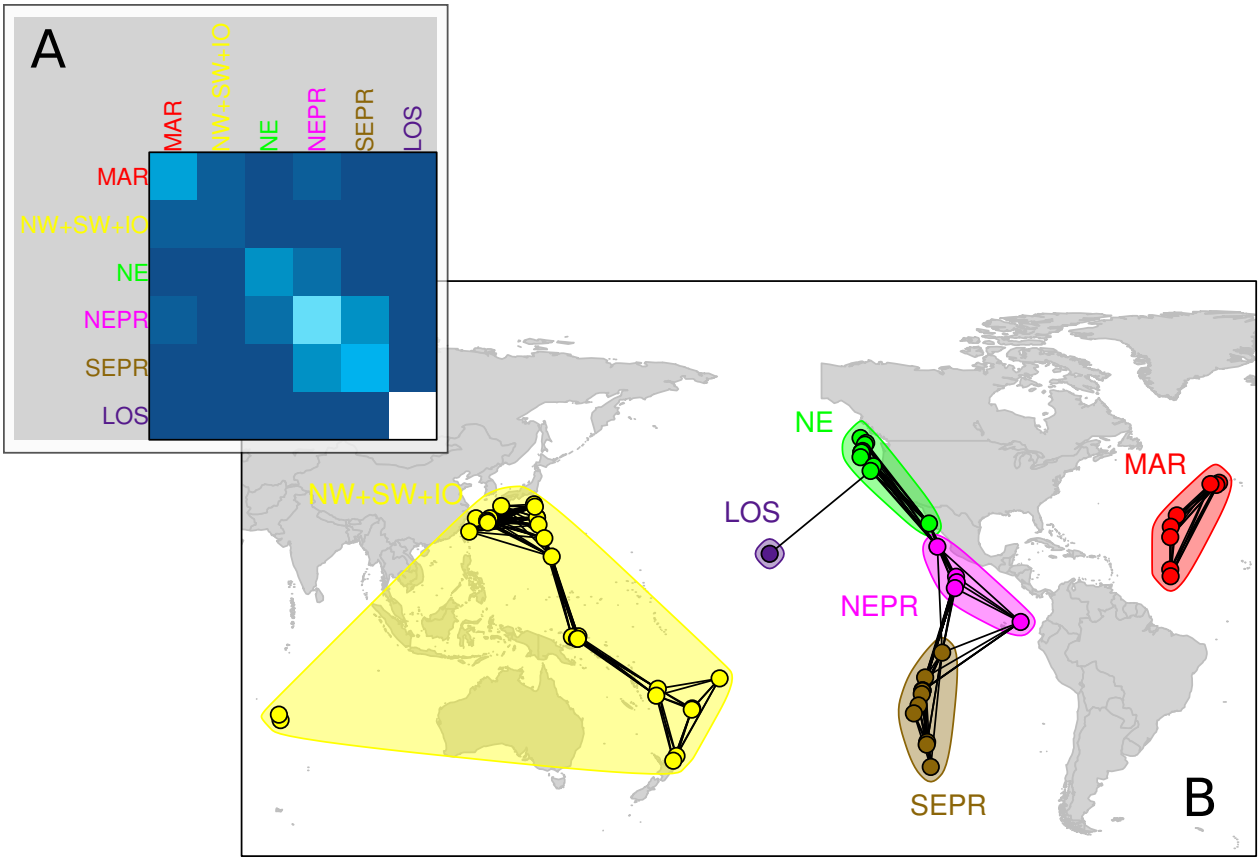

Figure 5: Spatial clustering of 63 hydrothermal fields in 6 groups based on similarities in species composition. (A) Heatmap showing average ecological distances within and between groups. (B) Map showing the structural network and the 6 groups: Garbage (GARB), NorthEast Pacific (NE), NorthEast Pacific Rise (NEPR), NorthEast Pacific Rise (SEPR), Mid-Atlantic Ridge (MAR) and NorthWest Pacific/Southwest Pacific/Indian Ocean (NW+SW+IO).